arXiv:1412.4098v3 [stat.ML] 6 Mar 2016

Manifold Matching using Shortest-Path Distance and Joint Neighborhood Selection

Cencheng Shen[a], Joshua T. Vogelstein[b], Carey E. Priebe[c,*]

[a]*Department of Statistics, Temple University, Philadelphia, PA 19122*
[b]*Department of Biomedical Engineering and Institute for Computational Medicine, Johns Hopkins University, Baltimore, MD 21218*
[c]*Department of Applied Mathematics and Statistics, Johns Hopkins University, Baltimore, MD 21218*

Abstract

We propose a nonlinear manifold matching algorithm to match multiple data sets using shortest-path distance and joint neighborhood selection. Based on the correspondence information, a neighborhood graph is jointly constructed; then the shortest-path distance within each data set is computed from the joint neighborhood graph, followed by embedding into and matching in a common low-dimensional Euclidean space. Our approach exhibits superior and robust performance for matching data from disparate sources, compared to algorithms that do not use shortest-path distance or joint neighborhood selection.

Keywords: manifold matching, shortest-path distance, joint neighborhood selection

1. Introduction

In the modern world it is becoming increasingly important to deal effectively with large amounts of high-dimensional data. For the purpose of data analysis, it is imperative to consider dimension reduction and embed the data into a low-dimensional space for subsequent analysis. Traditional linear embedding techniques have solid theoretical foundations and are widely used, e.g., principal component analysis (PCA) (Jolliffe, 2002; Tipping and Bishop, 1999) and multi-dimensional scaling (MDS) (Torgerson, 1952; Borg and Groenen, 2005; Cox and Cox, 2001) for a single data set, and canonical correlation analysis (CCA) (Hotelling, 1936; Bach and Jordan, 2005) for multiple data sets.

However, real data may exhibit nonlinear geometry, and discovering the underlying non-linearity can be beneficial for subsequent inference. Many manifold learning algorithms have been proposed to learn the intrinsic low-dimensional structure of nonlinear data, including Isomap (Tenenbaum et al., 2000; de Silva and Tenenbaum, 2003), locally linear embedding (LLE) (Saul and Roweis, 2000; Roweis and Saul, 2003), Hessian LLE (Donoho and Grimes, 2003), Laplacian eigenmaps (Belkin and Niyogi, 2003; He et al., 2005), local tangent space alignment (LTSA) (Zhang and Zha, 2004; Zhang et al., 2012), among many others. Most of them start with the assumption that the data are locally linear, explore the local geometry via the nearest-neighbor graph of the sample data, carry out transformation of the data based on the neighborhood graph, and eventually learn the low-dimensional manifold by optimizing some objective function. These nonlinear embedding algorithms usually serve as a preliminary feature extraction step enabling subsequent inference, and have achieved many successes in object recognition, image processing, etc.

In this paper, we consider the manifold match-

*Corresponding author
Email addresses: cshen@temple.edu (Cencheng Shen), jovo@jhu.edu (Joshua T. Vogelstein), cep@jhu.edu (Carey E. Priebe)

ing task for two or more data sets from disparate sources. Classical methods for identifying the relationship among multiple random variables are still very popular in theory and practice, such as canonical correlation (Hotelling, 1936; Kettenring, 1971; Hardoon et al., 2004) and Procrustes transformation (Sibson, 1978, 1979; Goldberg and Ritov, 2009; Gower and Dijksterhuis, 2004). But it has become a much more challenging task to match real data sets collected from distinct sources, such as the same document in different languages, an image and its descriptions, networks of the same actors on different social websites, etc.

There are many recent endeavors regarding data fusion and manifold matching (Lafon et al., 2006; Wang and Mahadevan, 2008; Wang et al., 2012; Sharma et al., 2012; Priebe et al., 2013). Similar to dimension reduction of a single data set, manifold matching can serve as a feature extraction step to explore multiple data sets, and has also been shown to help subsequent inference in object recognition, information retrieval, and transfer learning (Kim et al., 2007; Pan and Yang, 2010; Sun and Priebe, 2013; Sun et al., 2013; Shen et al., 2014). Furthermore, the matching task is important on its own and has been applied to explore multiple graphs and networks (Lyzinski et al., 2014; Vogelstein et al., 2015; Lyzinski et al., 2016): one such application is seeded graph matching, where two large networks are collected but only a percentage of training vertices have known correspondence; then the remaining vertices need to be matched properly to uncover the potential correspondence.

Due to the success of nonlinear embedding algorithms for a single data set, it seems natural that these algorithms can be combined into the matching framework to improve the matching performance when handling nonlinear data. A naive procedure is to pick one nonlinear algorithm, apply it to each data set separately, and match the embedded data sets. But such a naive procedure does not always guarantee a good matching performance, since many nonlinear embedding algorithms only preserve the local geometry up to some affine transformation (Goldberg and Ritov, 2008). And in the numerical experiments we will show that a direct matching of separate nonlinear embeddings can even deteriorate the matching performance when compared to linear embeddings.

To tackle the problem, in this paper we propose a manifold matching algorithm using shortest-path distance and joint neighborhood selection. By utilizing a robust distance measure that approximates the geodesic distance, and effectively combining the correspondence information into the embedding step, our proposed algorithm can significantly improve the matching quality from disparate data sources, compared to matching linear embeddings or separate nonlinear embeddings. Furthermore, the improvement in manifold matching motivates our follow-on work (Shen et al., 2016) regarding dependency discovery by local scale, which makes use of the joint neighborhood in a different manner.

The paper is organized as follows: In Section 2.1 we describe the matching task and the evaluation criteria in mathematical notations. The proposed nonlinear manifold matching algorithm is presented in Section 2.2, with various implementation issues discussed in Section 2.3. In Section 3 we illustrate the advantages of our algorithm via numerical simulations and real data experiments, using the simulated Swiss roll simulation and the Wikipedia document data sets with text and graph features. We summarize the results in Section 4. All codes and data are available on our website [1].

2. Manifold Matching

2.1. The Matching Framework

Suppose n objects are measured under two different sources. Then $X_l = \{x_{il}\} \in \Xi_l$ for $l = 1, 2$ are the actual data sets that are observed / collected, with $x_{i1} \sim x_{i2}$ for each i ($\sim$ means the two observations are matched in the context); alternatively, X_1 and X_2 can be thought as two different views of the same underlying data. This setting is extendable to more than two data sets, but for ease of presentation we mainly consider two data sets matching.

[1] `http://www.cis.jhu.edu/~cshen/`

Ξ_1 and Ξ_2 are potentially very different from each other, such as a flat manifold and its nonlinear transformation, an image and its description, or texts under different languages; and we assume $x_{il} \in \Xi_l$ is endowed with a distance measure Δ_l such that $\Delta_l(i, j) = dist(x_{il}, x_{jl})$. To match data sets from different spaces, we find two mappings $\rho_l : \Xi_l \to \mathbb{R}^d, l = 1, 2$ such that the mapped data $\hat{X}_l = \{\rho_l(x_{il})\}$ are matched into a common low-dimensional Euclidean space $\mathbb{R}^d$. A simple example of ρ_l can be MDS or PCA followed by CCA, which is a linear embedding and matching procedure.

Once the mappings are learned, for any new observations $y_1 \in \Xi_1$ and $y_2 \in \Xi_2$ of unknown correspondence, the learned mappings ρ_l can be applied to match the testing observations in the the low-dimensional Euclidean space, i.e., $\hat{y}_l = \rho_l(y_l) \in \mathbb{R}^d$. Ideally, a good matching procedure is able to correctly identify the correspondence of the new observations: if the testing observations are truly matched in the context, the mapped points should be very close to each other in the common Euclidean space; and if the testing observations are not matched, the mapped points should be far away from each other.

To evaluate a given matching algorithm, a natural criterion is the matching ratio used in seeded graph matching (Lyzinski et al., 2014): suppose sufficient training observations are given to learn the mappings, and there exist some testing observations of unknown correspondence in each space. Assume for each testing observation y_1 in Ξ_1, there is another testing observation $y_2 \in \Xi_2$ such that $y_1 \sim y_2$. Then they are correctly matched if and only if $\hat{y}_2$ is the nearest neighbor of $\hat{y}_1$ among all mapped testing data from Ξ_2. The matching ratio equals the percentage of correct matching of all testing observations, and a higher ratio indicates a better matching algorithm.

The matching ratio based on nearest neighbor is often conservative, and can be a very small number when matching disparate real data sets. In practice, it is often more interesting to consider all neighbors within a small threshold, or rank multiple neighbors up to a limit. To that end, the statistical testing power of the hypothesis $H_0 : y_1 \sim y_2$ considered in (Priebe et al., 2013) is another suitable criterion, which takes the Euclidean distance $\|\hat{y}_1 - \hat{y}_2\|$ as the test statistic. To estimate the testing power for given data, we first split all observations into matched training data pairs, matched testing data pairs, and unmatched testing data pairs; after learning ρ_l from the matched training data and applying them to all testing data, the test statistic under the null can be estimated from the matched testing pairs, and the test statistic under the alternative can be estimated from the unmatched testing pairs. Then the testing power at any type 1 error level is directly estimated from the empirical distributions of the test statistic, and a higher testing power indicates a better manifold matching algorithm.

We use both the testing power and the matching ratio for evaluation in the numerical experiments. Note that if the critical value at a given type 1 error level is used as a distance threshold, the testing power equals the percentage that the distance between the matched pair is no larger than the distance threshold. Because the matching ratio only considers the nearest neighbor of the matched pair, the testing power is always no smaller than the matching ratio.

2.2. Main Algorithm

The main algorithm is shown in Algorithm 1, which is referred to as MMSJ from now on. It is readily extendable to match more than two data sets; and some implementation details are discussed in the next subsection.

The algorithm can be summarized by a flowchart.

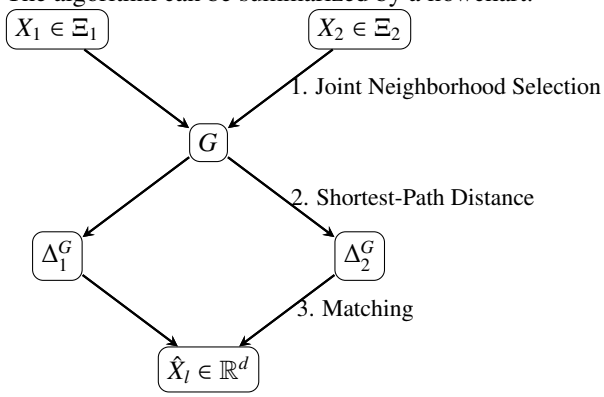

2.3. Discussions

In this subsection we discuss some implementation details of the nonlinear manifold matching algorithm.

Algorithm 1 Manifold Matching using Shortest-Path Distance and Joint Neighborhood Selection(MMSJ)

Input: Two data sets $\{X_l, l = 1, 2\}$ of sample size n with distance matrices Δ_l; a neighborhood parameter k, and a dimension choice d. We assume the two distance measures are properly scaled, e.g. both distance matrices are scaled to unit Frobenius norm.

Step 1: Construct an $n \times n$ graph adjacency G by k-nearest-neighbor on the joint distance measures $\sum_{l=1}^{2} \Delta_l$, i.e., $G(i, j) = 1$ if and only if $\sum_l \Delta_l(x_{il}, x_{jl})$ is among the smallest k elements in the set $\{\sum_l \Delta_l(x_{il}, x_{ql}), q = 1, \ldots, n\}$.

Step 2: For each data set X_l, calculate the corresponding shortest-path distance matrix Δ_l^G from the graph G: for each i, j, first initialize

$$\Delta_1^G(i, j) = \begin{cases} \Delta_1(i, j), & \text{if } G(i, j) = 1, \\ \infty, & \text{otherwise ;} \end{cases}$$

then iterate through $q = 1, \ldots, n$ and set

$$\Delta_1^G(i, j) = \min\{\Delta_1^G(i, j), \Delta_1^G(i, q) + \Delta_1^G(q, j)\};$$

the final Δ_1^G is the shortest-path distance matrix of X_1 by joint neighborhood. Repeat the above for Δ_2 using the same G to yield Δ_2^G.

Step 3: Embed the shortest-path distance matrices Δ_l^G into $\mathbb{R}^d$ by MDS, and use the Procrustes matching to yield the matched data $\hat{X}_l$, i.e., denote $\tilde{X}_l$ as the embedded data by MDS, the Procrustes matching finds a $d \times d$ rotation matrix by

$$P = \arg\min_{P'P=I} \|P\tilde{X}_1 - \tilde{X}_2\|_F^2,$$

then sets $\hat{X}_1 = P\tilde{X}_1$ and $\hat{X}_2 = \tilde{X}_2$.

Output: Two mapped data sets $\hat{X}_l \in \mathbb{R}^{d \times n}$ for $l = 1, 2$; and the learned Procrustes transformation P.

In step 1, a joint neighborhood graph is constructed from the distance matrices of both data sets, which ensures consistent neighborhood selection in case of noisy or nonlinear data sets, and are better than separate neighborhood for matching. Note that it is necessary for the distance matrices to be properly scaled so that the joint neighborhood selection is meaningful, or one may use a weighted sum of distances or rank-based method to derive a joint neighborhood instead.

In step 2, the shortest-path distance within each data set is calculated using the joint neighborhood graph, which can recover the geodesic distance of isometric manifolds with high probability under certain sampling condition (Bernstein et al., 2000; de Silva and Tenenbaum, 2003). Conceptually, the shortest-path distance keeps intact the distances within the local neighborhood while enlarging far away distances. Computationally, the shortest-path distance matrix can be effectively implemented by Floyd's algorithm or Dijkstra's algorithm (Tenenbaum et al., 2000), which can be further sped up by choosing a small set of landmark points (de Silva and Tenenbaum, 2003; Bengio et al., 2003). In comparison, many other nonlinear algorithms like LLE, LTSA, and Laplacian eigenmaps involve a normalization step and only preserve the local geometry up to some affine transformation (Goldberg and Ritov, 2008), such that they may fail to match noisy or disparate data sets.

In step 3, the shortest-path distance matrices are embedded by MDS into $\mathbb{R}^d$ followed by the Procrustes matching. Alternatively, the CCA matching and the joint MDS matching described in (Priebe et al., 2013) can also be used, which are implemented in our code as well. They have their pros and cons in linear matching as investigated in (Priebe et al., 2013; Fishkind et al., 2016); but the advantage of the MMSJ algorithm mostly lies in step 1 and 2 rather than the traditional step 3, as MMSJ always exhibits significant improvement no matter the matching method in step 3. Thus in the paper we mainly use the Procrustes matching for ease of presentation. Once the mappings are learned from the training data, they can be applied to match any new testing observations.

Other than MMSJ, the naive procedure to embed and match by a given linear/nonlinear embedding algorithm (MDS / Isomap / LLE / LTSA, etc.) is as follows: X_1 and X_2 are separately embedded from the original space (using the original data / original distance matrices) to $\mathbb{R}^d$ by the specific algorithm, and the embedded data are then matched by the Procrustes matching. Note that MDS / Isomap / LLE can all operate directly on a distance matrix, but some nonlinear algorithms like LTSA have to start with the Euclidean data rather than a distance measure; so if only the distance matrices are available, MDS is first used to embed the distance matrices into a Euclidean space $\mathbb{R}^{d'}$ with $d' \geq d$, followed by LTSA to embed into $\mathbb{R}^d$ then Procrustes matching.

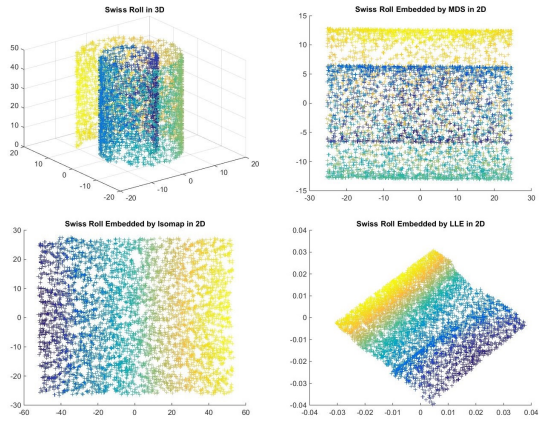

Figure 1: The 3D Swiss roll data set (left top), its 2D embedded data by MDS (right top), 2D embedding by Isomap at neighborhood choice $k = 10$ (left bottom), and 2D embedding by LLE at $k = 10$ (right bottom).

3. Numerical Experiments

In this section we demonstrate the numerical advantages of the proposed manifold matching algorithm, with MDS, Isomap, LLE, and LTSA as the benchmarks. Overall, we observe that our algorithm is significantly better than all the benchmarks in matching ratio and testing power.

3.1. Swiss Roll Simulation

The Swiss roll data from (Tenenbaum et al., 2000) is a 3D data set representing a nonlinear manifold, but intrinsically generated by points on a 2D linear manifold. Figure 1 shows the 3D Swiss roll data with 5000 points in colors, along with its 2D embeddings by MDS, Isomap and LLE. Clearly MDS fails to recognize the nonlinear geometry while both Isomap and LLE succeed. However, the LLE embedding has a distorted geometry, while the Isomap embedding is similar to the underlying 2D linear manifold.

For the first simulation, we match the 3D Swiss roll with its underlying 2D linear manifold. A total of $n = 1000$ points from the 3D Swiss roll are randomly generated as the first data set X_1, and the corresponding points on the underlying 2D linear manifold are taken as the second data set X_2. Thus X_1 and X_2 are matched training data with distinct geometries. Once the training data are matched, we embed and apply the learned mappings to new testing observations y_1 and y_2 in each space.

We set the parameters as $k = 10$ and $d = 2$, and use 100 matched testing pairs and 100 unmatched testing pairs to calculate the test statistic under the null and the alternative. Repeat the above for 100 Monte-Carlo replicates, the mean testing powers with respect to the increasing type 1 error level is plotted in Figure 2(a). The matching ratio is estimated based on the same 100 matched testing pairs from 100 MC replicates, which is shown in the first row of Table 1. The proposed MMSJ is clearly much better than all the benchmarks in both the testing power and the matching ratio; MDS and LTSA have the worst matching performance, while Isomap and LLE perform reasonably good.

Next we check the robustness of the manifold matching algorithms against noise, by adding white noise to the linear data set X_2. The noise is independently identically distributed as $N(0, \epsilon I_{2\times 2})$, and the same testing procedure is applied to estimate the testing powers at each noise level. The powers are plotted in Figure 2(b) with respect to the increasing noise level $\epsilon = 0, 1, 2, \ldots, 10$ at type 1 error level 0.05. The proposed MMSJ algorithm is still better

than all the benchmarks as the noise level increases.

For the third simulation, we match two linear data sets: instead of the 3D Swiss roll, we use the 2D LLE embedding as X_1; thus X_1 and X_2 are both linear with some differences. Using the same procedure and parameters as the first simulation, the testing power is plotted in Figure 2(d) and the matching ratio is shown in the second row of Table 1. In this case MMSJ, MDS, and Isomap have similar performances with LLE and LTSA being slightly worse, which indicates that the shortest-path distance is more robust than nonlinear algorithms like LLE and LTSA as pointed out in Section 2.3. Note that the matching ratio and testing power here are lower than those in the first simulation, because the LLE embedding has a distorted linear geometry.

3.2. Wikipedia Articles Experiments

In the real data experiments we apply the manifold matching algorithm to match disparate features of Wikipedia articles. The raw data contains 1382 pairs of articles from Wikipedia English and its corresponding French translations, within the 2-neighborhood of the English article "Algebraic Geometry". On Wikipedia, the same articles of different languages are almost never the exact translations of each other, because they are very likely written by different people and their contents may differ in many ways.

For the English articles and their French translations, a text feature and a graph feature are collected separately under each language: for the texts of each article, we use latent semantic indexing (LSI) (Deerwester et al., 1990) followed by cosine dissimilarity to construct two dissimilarity matrices TE and TF (standing for English texts and French texts); for the networks, two shortest-path distance matrices GE and GF (standing for English graph and French graph) are calculated based on the Internet hyperlinks of the articles under each language setting, with any path distance larger than 4 imputed to be 6 to avoid infinite distances and scaling issues.

Therefore, there exist four different data representations for the same Wikipedia articles, making TE, TF, GE, GF matched in the context. Furthermore, as the text matrices are derived by cosine similarity

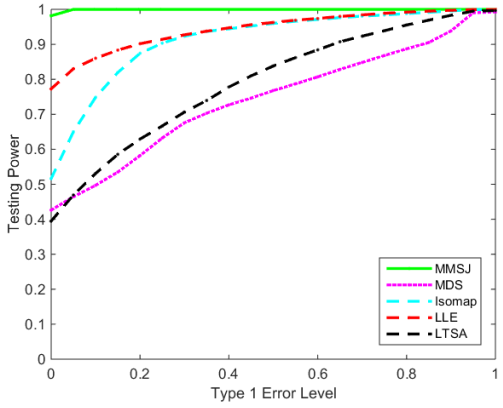

(a)

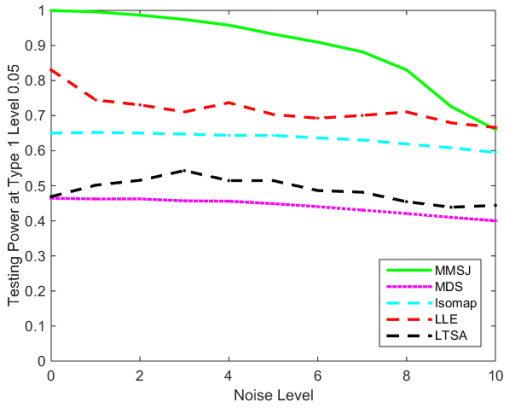

(b)

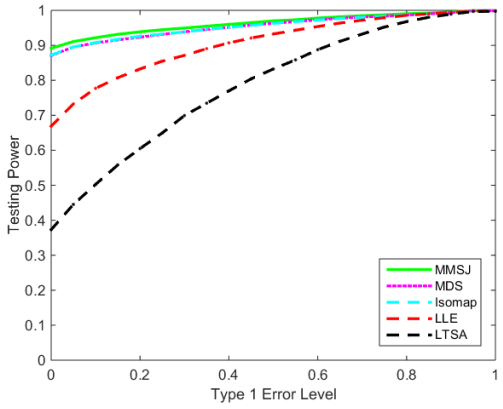

(c)

Figure 2: (a): Testing Power of 3D Swiss Roll versus its 2D Underlying Linear Manifold; (b): Testing Power of 3D Swiss Roll versus its 2D Underlying Linear Manifold with Increasing Noise; (c): Testing Power of 2D LLE Embedding of Swiss Roll versus its 2D Underlying Linear Manifold.

Table 1: Swiss Roll Matching Ratio

Data Pair	MMSJ	MDS	Isomap	LLE	LTSA
(3D Swiss Roll, 2D Linear Manifold)	**0.9787**	0.0200	0.1324	0.2123	0.0697
(2D LLE Embedding, 2D Linear Manifold)	**0.1771**	0.1386	0.1412	0.1248	0.0828

while the graph matrices are based on the shortest-path distance, the former probably have nonlinear geometries while the latter should be close to linear.

For each Monte-Carlo replicates, we randomly pick $n = 500$ pairs of training observations, 100 pairs of testing matched observations and 100 pairs of testing unmatched observations for evaluation. The parameters are set as $k = 20$, $d = 10$, $d' = 50$ (for LTSA only), and the manifold matching algorithms are applied for every two data sets matching combination. After 100 Monte-Carlo replicates, the mean matching ratio is reported in Table 2, the estimated testing power is presented in Table 3 at type 1 error level 0.05, and the power curves for some matching combinations are plotted in Figure 3.

Clearly MMSJ achieves the best performance throughout all combinations. From the tables and figures, we further observe that without using shortest-path distance or joint neighborhood, separate nonlinear embeddings from LLE or LTSA are worse than the linear MDS embeddings in matching; Isomap does fairly well in the testing power as it also uses shortest-path distance, but can still be similar or slightly inferior to MDS in the matching ratio occasionally; and our proposed MMSJ algorithm is consistently the best manifold matching algorithm in both the testing power and the matching ratio throughout.

We should point out that the matching performance depends on the parameters k and d, and the testing power / matching ratio of each algorithm can be slightly different by changing the parameters. For example, the MDS testing power is 0.49 for matching (TE, TF) at $d = 10$, which can be improved to 0.7 by varying d; MMSJ has the best power of 0.82, which can be increased to 0.9 by cross-validating the parameters as well. So far we have used fixed parameter choices to offer meaningful comparisons throughout all algorithms and matching combinations, but our MMSJ algorithm is in fact robust against parameter choices. As an example, in Figure 4 we show the MMSJ and Isomap testing powers for matching (TE, GE) (the best two algorithms in our matching experiments) against different choices of d and k, for which d ranges from 2 to 30 and k ranges from 10 to 30. It is clear that MMSJ is always better than Isomap in matching and attains close-to-optimal testing power in a large range of parameter choices; furthermore, the best MMSJ testing power is 0.55 while the best Isomap testing power is 0.45. The same robustness holds for MMSJ under all other matching combinations.

Table 2: Wikipedia Features Matching Ratio

Data Pair	MMSJ	MDS	Isomap	LLE	LTSA
(TE, TF)	**0.2942**	0.2546	0.2003	0.1265	0.0491
(TE, GE)	**0.1209**	0.0675	0.0866	0.0143	0.0260
(TF, GF)	**0.0624**	0.0419	0.0522	0.0134	0.0144
(GE, GF)	**0.1347**	0.1280	0.1081	0.0157	0.236
(TE, GF)	**0.0677**	0.0429	0.0560	0.0132	0.0138
(TF, GE)	**0.0946**	0.0545	0.0698	0.0132	0.0238

Table 3: Wikipedia Features Testing Power at Type 1 Error Level 0.05

Data Pair	MMSJ	MDS	Isomap	LLE	LTSA
(TE, TF)	**0.8124**	0.4974	0.7476	0.3594	0.1930
(TE, GE)	**0.5184**	0.2563	0.4255	0.0948	0.1116
(TF, GF)	**0.2782**	0.1128	0.1877	0.0903	0.1028
(GE, GF)	**0.3108**	0.2141	0.2485	0.0961	0.1063
(TE, GF)	**0.3199**	0.1130	0.2141	0.0923	0.1021
(TF, GE)	**0.4464**	0.2114	0.3595	0.0943	0.1064

4. Concluding Remarks

In summary, we propose a nonlinear manifold matching algorithm using shortest-path distance and

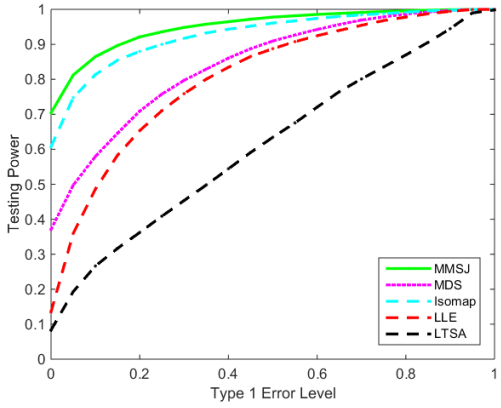

(a)

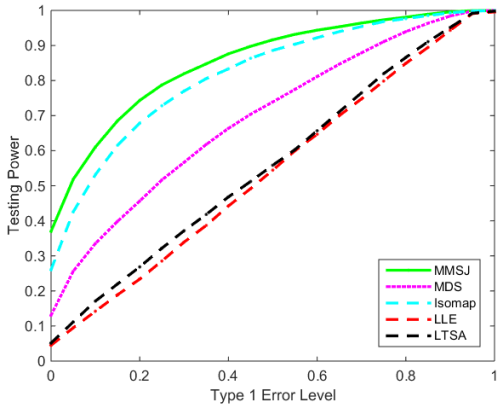

(b)

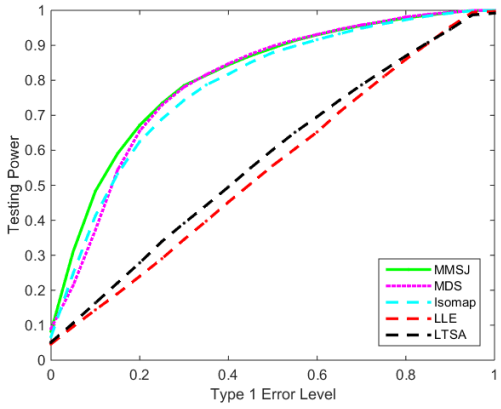

(c)

Figure 3: (a): Testing Power of Wikipedia English Text versus French Text. (b): Testing Power of Wikipedia English Text versus English Text. (c): Testing Power of Wikipedia English Graph versus French Graph.

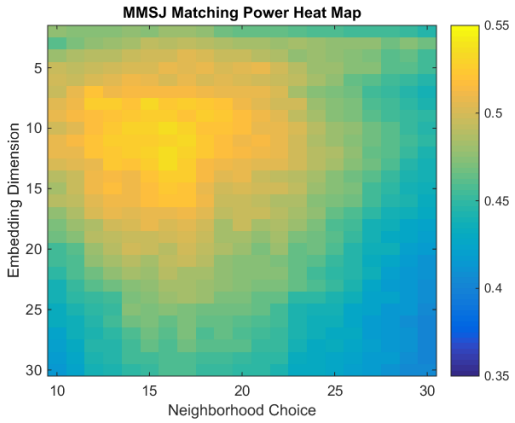

(a)

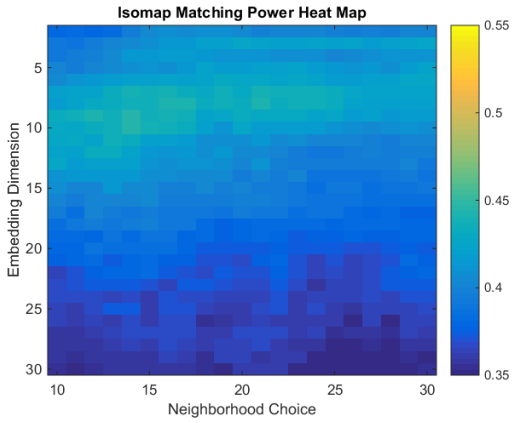

(b)

Figure 4: Testing Power of Wikipedia English Text and English Graph with respect to Different Dimension Choices and Neighborhood Sizes at Type 1 Error Level 0.05.

joint neighborhood. The algorithm is straightforward to implement, and achieves superior and robust performance. It is able to significantly improve the testing power and matching ratio when matching data of distinct geometries, and is robust against noise and model selection. Our experiments indicate that the shortest-path distance and joint neighborhood selection are two key catalysts behind the improvement of the matching performance.

Acknowledgment

This work is partially supported by National Security Science and Engineering Faculty Fellowship (NSSEFF), Johns Hopkins University Human Language Technology Center of Excellence (JHU HLT COE), and the XDATA program of the Defense Advanced Research Projects Agency (DARPA) administered through Air Force Research Laboratory contract FA8750-12-2-0303.

This work is also supported by the Defense Advanced Research Projects Agency (DARPA) SIMPLEX program through SPAWAR contract N66001-15-C-4041 and DARPA GRAPHS N66001-14-1-4028.

References

Bach, F.R., Jordan, M.I., 2005. A Probabilistic Interpretation of Canonical Correlation Analysis. Technical Report. Department of Statistics, UC Berkeley.

Belkin, M., Niyogi, P., 2003. Laplacian eigenmaps for dimensionality reduction and data representation. Neural Computation 15, 1373–1396.

Bengio, Y., Paiement, J.F., Vincent, P., 2003. Out-of-sample extensions for LLE, Isomap, MDS, Eigenmaps, and Spectral Clustering, in: Advances in Neural Information Processing Systems, MIT Press. pp. 177–184.

Bernstein, M., de Silva, V., Langford, J.C., Tenenbaum, J.B., 2000. Graph approximations to geodesics on embedded manifolds.

Borg, I., Groenen, P., 2005. Modern Multidimensional Scaling: theory and applications. Springer-Verlag.

Cox, T., Cox, M., 2001. Multidimensional Scaling. Chapman and Hall.

Deerwester, S., Dumais, S., Landauer, T., Furnas, G., Harshman, R., 1990. Indexing by latent semantic analysis. Journal of the American Society of Information Science 41, 391–407.

Donoho, D., Grimes, C., 2003. Hessian eigenmaps: New locally linear embedding techniques for high-dimensional data, in: Proceedings of the National Academy of Arts and Sciences, pp. 5591–5596.

Fishkind, D., Shen, C., Park, Y., Priebe, C.E., 2016. On the incommensurability phenomenon. Journal of Classification accepted.

Goldberg, Y., Ritov, Y., 2008. Manifold learning: the price of normalization. Journal of Machine learning research 9, 1909–1939.

Goldberg, Y., Ritov, Y., 2009. Local Procrustes for manifold embedding: a measure of embedding quality and embedding algorithms. Machine learning 77, 1–25.

Gower, J.C., Dijksterhuis, G.B., 2004. Procrustes Problems. Oxford University Press.

Hardoon, D.R., Szedmak, S., Shawe-Taylor, J., 2004. Canonical correlation analysis: An overview with application to learning methods. Neural Computation 16, 2639–2664.

He, X., Yan, S., Hu, Y., Niyogi, P., Zhang, H., 2005. Face recognition using Laplacianfaces. IEEE Transactions on Pattern Analysis and Machine Intelligence 27, 328–340.

Hotelling, H., 1936. Relations between two sets of variates. Biometrika 28, 321–377.

Jolliffe, I.T., 2002. Principal Component Analysis. 2nd ed., Springer.

Kettenring, J.R., 1971. Canonical analysis of several sets of variables. Biometrika 58, 433–451.

Kim, T.K., Kittler, J., Cipolla, R., 2007. Discriminative learning and recognition of image set classes using canonical correlations. IEEE Transactions on Pattern Analysis and Machine Intelligence 29, 1005–1018.

Lafon, S., Keller, Y., Coifman, R., 2006. Data fusion and multi-cue data matching by difusion maps. IEEE transactions on Pattern Analysis and Machine Intelligence 28, 1784–1797.

Lyzinski, V., Fishkind, D., Fiori, M., Vogelstein, J.T., Priebe, C.E., Sapiro, G., 2016. Graph matching: Relax at your own risk. IEEE Transactions on Pattern Analysis and Machine Intelligence 38, 60–73.

Lyzinski, V., Fishkind, D., Priebe, C.E., 2014. Seeded graph matching for correlated Erdos-Renyi graphs. Journal of Machine Learning Research 15, 3513–3540.

Pan, S.J., Yang, Q., 2010. A survey on transfer learning. IEEE Transactions on Knowledge and Data Engineering 22, 1345–1359.

Priebe, C.E., Marchette, D.J., Ma, Z., Adali, S., 2013. Manifold matching: Joint optimization of fidelity and commensurability. Brazilian Journal of Probability and Statistics 27, 377–400.

Roweis, S.T., Saul, L.K., 2003. Think globally, fit locally: Unsupervised learning of low dimensional manifolds. Journal of Machine Learning Research 4, 119–155.

Saul, L.K., Roweis, S.T., 2000. Nonlinear dimensionality reduction by locally linear embedding. Science 290, 2323–2326.

Sharma, A., Kumar, A., III, H.D., Jacobs, D., 2012. Generalized multiview analysis: A discriminative latent space, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

Shen, C., Priebe, C.E., Maggioni, M., Vogelstein, J.T., 2016. Dependence discovery from multimodal data via multiscale graph correlation. Submitted .

Shen, C., Sun, M., Tang, M., Priebe, C.E., 2014. Generalized canonical correlation analysis for classification. Journal of Multivariate Analysis 130, 310–322.

Sibson, R., 1978. Studies in the robustness of multidimensional scaling: Procrustes statistics. Journal of the Royal Statistical Society. Series B 40, 234–238.

Sibson, R., 1979. Studies in the robustness of multidimensional scaling: Perturbation analysis of classical scaling. Journal of the Royal Statistical Society. Series B 41, 217–229.

de Silva, V., Tenenbaum, J.B., 2003. Global versus local methods in nonlinear dimensionality reduction. Advances in Neural Informaiton Processing Systems 15, 721–728.

Sun, M., Priebe, C.E., 2013. Efficiency investigation of manifold matching for text document classification. Pattern Recognition Letters 34, 1263–1269.

Sun, M., Priebe, C.E., Tang, M., 2013. Generalized canonical correlation analysis for disparate data fusion. Pattern Recognition Letters 34, 194–200.

Tenenbaum, J.B., de Silva, V., Langford, J.C., 2000. A global geometric framework for nonlinear dimension reduction. Science 290, 2319–2323.

Tipping, M.E., Bishop, C.M., 1999. Probabilistic principal component analysis. Journal of the Royal Statistical Society, Series B 61, 611–622.

Torgerson, W., 1952. Multidimensional Scaling: I. Theory and method. Psychometrika.

Vogelstein, J., Conroy, J., Lyzinski, V., Podrazik, L., Kratzer, S., Harley, E., Fishkind, D., Vogelstein, R., Priebe, C., 2015. Fast approximate quadratic programming for graph matching. PLOS ONE 10, e0121002.

Wang, C., Liu, B., Vu, H., Mahadevan, S., 2012. Sparse manifold alignment, in: Technical Report, UMass Computer Science UM-2012-030.

Wang, C., Mahadevan, S., 2008. Manifold alignment using Procrustes analysis, in: Proceedings of the 25th International Conference on Machine Learning.

Zhang, Z., Wang, J., Zha, H., 2012. Adaptive manifold learning. IEEE Transactions on Pattern Analysis and Machine Intelligence 34, 253–265.

Zhang, Z., Zha, H., 2004. Principal manifolds and nonlinear dimensionality reduction via tangent space alignment. SIAM Journal on Scientific Computing 26, 313–338.